

Adaptive tracking of representative cycles in regular and zigzag persistent homology

Jennifer Gamble^{*}, Harish Chintakunta^{**}, and Hamid Krim^{**}

^{*}Electrical and Computer Engineering, North Carolina State University

^{**}Coordinated Science Laboratory, University of Illinois Urbana Champaign

Abstract

Persistent homology and zigzag persistent homology are techniques which track the homology over a sequence of spaces, outputting a set of intervals corresponding to birth and death times of homological features in the sequence. This paper presents a method for choosing a homology class to correspond to each of the intervals at each time point. For each homology class a specific representative cycle is stored, with the choice of homology class and representative cycle being both geometrically relevant and compatible with the birth-death interval decomposition. After describing the method in detail and proving its correctness, we illustrate the utility of the method by applying it to the study of coverage holes in time-varying sensor networks.

1 Introduction

The field of topological data analysis [2] has been blossoming in recent years, and many more statisticians, computer scientists and engineers are beginning to use topological tools to study their data. The most popular and successful of these tools is persistent homology, a method which characterizes a space (or object) by using a multi-scale description of its topological features. These include things like the number of connected components, holes, or voids. A variant of regular persistent homology is zigzag persistent homology, which describes the topological features as they change over a sequence of spaces.

In this paper, we propose an algorithm for obtaining specific representative cycles to track homological features over a sequence of simplicial complexes in a geometrically meaningful way. We include some explorations where this method is applied to track coverage holes in time-varying sensor networks, with fruitful results.

In Section 2, we build up the foundational terminology and notations required for our discussions. This includes brief descriptions of simplicial homology and zigzag persistence. In Section 3 we describe our method, and give an algorithm for implementing it on a sequence of simplicial complexes. Section 4 proves the correctness of our algorithm, and Section 5 illustrates the utility of the method for analysis of coverage properties of time-varying sensor networks.

^{*}jpgamble@ncsu.edu (Jennifer Gamble, corresponding author), ^{**}hkchinta@ncsu.edu (Harish Chintakunta), ^{**}ahk@ncsu.edu (Hamid Krim).

2 Terminology and notation

2.1 Simplicial complexes and homology

We will use ideas from simplicial homology theory throughout, so some basic notations are introduced here. For a general reference on algebraic topology (including simplicial homology), see [8]. We define the notation for simplicial complexes, homology groups and classes.

A k -simplex $\sigma = [v_0, v_1, \dots, v_k]$ is a set of $k + 1$ vertices. Any subset of the $k + 1$ vertices forming a simplex is called a *face* of the simplex, where each face is also a simplex itself. A *simplicial complex*, K , is a set of simplices such that any simplex in K also has all of its faces in K , and the intersection of any two simplices σ_1 and σ_2 in K is a face of both σ_1 and σ_2 .

Given a simplicial complex K , the *chain spaces* $\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \dots$, are vector spaces, where \mathcal{C}_k has the k -simplices as basis elements. The specific structure of the simplicial complex is encoded in the *boundary maps* $\partial_1, \partial_2, \dots$. For k -simplex $\sigma = [v_0, v_1, \dots, v_k]$, the boundary map $\partial_k : \mathcal{C}_k \rightarrow \mathcal{C}_{k-1}$ maps σ onto the alternating sum of its faces:

$$\partial_k \sigma = \sum_{i=0}^k (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_k]$$

where \hat{v}_i indicates vertex v_i is removed. Then the k -th homology group is

$$H_k(K) = \ker(\partial_k) / \text{im}(\partial_{k+1})$$

and the k -th Betti number (denoted β_k) of the simplicial complex K is the rank of $H_k(K)$.

An element c in the chain space \mathcal{C}_k is a linear combination of k -simplices $\sigma_1, \dots, \sigma_{n_k}$,

$$c = \sum_{i=1}^{n_k} a_i \sigma_i$$

The coefficients a_i come from a field \mathbb{F} (such as the real numbers). We choose to perform our computations over the field $\mathbb{Z}_2 = \{0, 1\}$. Any chain with boundary zero (i.e. any c such that $\partial_k c = 0$) is called a *cycle*, and so $\ker(\partial_k)$ is the set of all k -cycles. All boundaries of simplices are themselves cycles (i.e. $\text{im}(\partial_{k+1}) \subseteq \ker(\partial_k)$), so $\partial_k \partial_{k+1} = 0$, and homology corresponds to “cycles which are not boundaries”. Two cycles c_1 and c_2 are *homologous* (written $c_1 \sim c_2$) if their difference can be written as a linear combination of boundaries. The set of all cycles that are homologous to a given cycle (say c) is called a *homology class* (denoted $[c]$). When a specific cycle is chosen to represent an entire homology class, it is called a *representative cycle*.

When writing an equation about the homology of a space, we use a general notation of $H(K)$, which can be taken to mean that the total homology $H_*(K)$, or a specific $H_p(K)$, could be inserted into the equation in the place of $H(K)$. Similarly, we use $\beta(K) = \text{rank}(H(K))$ as a general notation for the associated Betti number. For the later applications to sensor networks, and for visualization purposes, it is convenient to think of $H(K)$ to mean $H_1(K)$.

2.2 Zigzag persistence

The theory of zigzag persistent homology is concerned with how the homology changes over a sequence of spaces. There are mathematical results [4] showing that the changing homology of such a sequence can be expressed uniquely in terms of birth and death times

Case	$V_i \xleftrightarrow{p_i} V_{i+1}$
1. Birth by addition	$H_d(K_i) \xrightarrow{f_i} H_d(K_{i+1})$
2. Birth by removal	$H_{d-1}(K_i) \xleftarrow{g_i} H_{d-1}(K_{i+1})$
3. Death by addition	$H_d(K_i) \xrightarrow{f_i} H_d(K_{i+1})$
4. Death by removal	$H_{d-1}(K_i) \xleftarrow{g_i} H_{d-1}(K_{i+1})$

Table 1: Indicates the relationship between the dimension d of simplex σ being added or removed, and the dimension of homology space which is affected.

of homological features in the sequence. There is also an algorithm [3] for computing this birth-death decomposition for a given sequence of simplicial complexes.

Consider a sequence of simplicial complexes K_1, K_2, \dots, K_n , connected by either forward inclusion maps $K_i \rightarrow K_{i+1}$ or backward inclusion maps $K_i \leftarrow K_{i+1}$. We write this sequence as

$$K_1 \longleftrightarrow K_2 \longleftrightarrow \dots \longleftrightarrow K_n$$

The inclusion maps induce linear maps between the associated homology spaces $V_i = H(K_i)$, which we write as the *zigzag persistence module*

$$\mathbb{V} = V_1 \xleftrightarrow{p_1} V_2 \xleftrightarrow{p_2} \dots \xleftrightarrow{p_{n-1}} V_n \xrightarrow{p_n} V_{n+1} \quad (1)$$

where $K_i \rightarrow K_{i+1}$ induces the forward map $V_i \xrightarrow{f_i} V_{i+1}$, and $K_i \leftarrow K_{i+1}$ induces the backward map $V_i \xleftarrow{g_i} V_{i+1}$. Regardless of the direction, we use $i(\cdot)$ to denote the inclusion map between consecutive simplicial complexes. We further assume that consecutive simplicial complexes K_i and K_{i+1} differ by exactly one simplex, so $K_{i+1} = K_i \cup \{\sigma\}$ (in the forward case), or $K_{i+1} = K_i - \{\sigma\}$ (in the backward case).

When σ is a d -simplex, its addition results in either an increase in the dimension of the d -dimensional homology space, or a decrease in the dimension of the $(d-1)$ -dimensional homology space. Similarly, the removal of a d -simplex σ results in either an increase in the $(d-1)$ -dimensional homology, or a decrease in the d -dimensional homology. When the dimension of the homology space increases, we refer to this as a *birth*, and when the dimension decreases, we refer to this as a *death*.

$$\text{Birth} : \dim(V_{i+1}) = \dim(V_i) + 1$$

$$\text{Death} : \dim(V_{i+1}) = \dim(V_i) - 1$$

Each inclusion between simplicial complexes will induce maps between the homology spaces of all dimensions, but these maps will be simple identity maps in all dimensions except for one. This will depend on whether the addition or removal of d -simplex σ results in a birth or a death. For the addition or removal of a d -simplex σ , the map $V_i \xleftrightarrow{p_i} V_{i+1}$ on the corresponding homology zigzag module, will be interpreted as a forward or backward linear map between the appropriate-dimensional homology spaces, as summarized in Table 1.

A main result from the theory of zigzag persistence, is that a zigzag module such as in Equation 1 has an interval decomposition,

$$\mathbb{V} \cong \mathbb{I}(b_1, d_1) \otimes \mathbb{I}(b_2, d_2) \otimes \dots \otimes \mathbb{I}(b_m, d_m) \quad (2)$$

which is unique up to isomorphism, and is equivalently expressed as the multiset of pairings of births and deaths in the sequence, and represented as integer intervals, called the *zigzag persistence* of \mathbb{V} [4]

$$\text{Pers}(\mathbb{V}) = \{[b_j, d_j] \mid j = 1, \dots, m\}. \quad (3)$$

These are interpreted as birth and death times of homological features in the sequence.

2.3 Right filtration

Given a zigzag module \mathbb{V} as in Equation 1, the zigzag persistence algorithm [3] computes the interval decomposition in Equation 2 by keeping track of a right filtration $R(\mathbb{V})$ on the spaces. The right filtration $R(\mathbb{V})$ is computed incrementally, and results in a filtration (a nested sequence of subspaces) on V_n , along with a birth time associated to each quotient space. A right filtration on V_i is denoted

$$\mathcal{R}_i = (R_i^0, R_i^1, \dots, R_i^i) \quad (4)$$

where $R_i^0 \leq R_i^1 \leq \dots \leq R_i^i$ and $R_i^i = V_i$. The quotients $R_i^1/R_i^0, R_i^2/R_i^1, \dots, R_i^i/R_i^{i-1}$ are each associated with a birth time b_i^j (for $j = 0, \dots, i$), which are recorded in the vector

$$\mathbf{b}_i = (b_i^1, b_i^2, \dots, b_i^i) \quad (5)$$

We may write the quotients as

$$\mathcal{R}'_i = (R_i^1/R_i^0, R_i^2/R_i^1, \dots, R_i^i/R_i^{i-1})$$

The computation of a right filtration is defined inductively, depending on whether the map from V_i to V_{i+1} is a forward map $\xrightarrow{f_i}$ or a backward map $\xleftarrow{g_i}$. For a single vector space V_1 , we have the base case of $i = 1$, and we define

$$\mathcal{R}_1 = (\mathbf{0}, V_1) \text{ and } \mathbf{b}_1 = (0)$$

In the inductive step, if we are given \mathcal{R}_i and \mathbf{b}_i as in Equations 4 and 5 above, then

- If $V_i \xrightarrow{f_i} V_{i+1}$, then

$$\begin{aligned} \mathcal{R}_{i+1} &= (f_i(R_i^0), f_i(R_i^1), \dots, f_i(R_i^i), V_{i+1}) \\ \mathbf{b}_{i+1} &= (b_i^1, b_i^2, \dots, b_i^i, i+1) \end{aligned} \quad (6)$$

- If $V_i \xleftarrow{g_i} V_{i+1}$, then

$$\begin{aligned} \mathcal{R}_{i+1} &= (\mathbf{0}, g_i^{-1}(R_i^0), g_i^{-1}(R_i^1), \dots, g_i^{-1}(R_i^i)) \\ \mathbf{b}_{i+1} &= (i+1, b_i^1, b_i^2, \dots, b_i^i) \end{aligned} \quad (7)$$

Since we assume that consecutive simplicial complexes differ by at most one simplex, the change in dimension between V_i and V_{i+1} is at most 1. Similarly, the dimension of the quotient space R_i/R_{i+1} is either 0 or 1, for $i = 1, \dots, n$, with their total dimension equaling that of V_i . The dimension of V_i is the rank of the homology group for K_i (the Betti number, $\beta(K_i)$), which is at most i :

$$\dim(V_i) = \text{rank}(\text{H}(K_i)) = \beta(K_i) \leq i$$

For example, the dimension of the quotient spaces will be a sequence of 0's and 1's

$$\dim(R_i^1/R_i^0, R_i^2/R_i^1, \dots, R_i^i/R_i^{i-1}) = (0, 0, 1, 1, 0, \dots, 1, 0)$$

Note that choosing one homology class from each of the nonzero quotient spaces results in a basis for V_i . The right filtration on V_i can then be described using the unique subspaces in the right filtration (which have corresponding quotient spaces of dimension 1). Indexing the nonzero quotient spaces by $j_1, \dots, j_{\beta(K_i)}$, define $W_i^k = R_i^{j_k}$ for the spaces $R_i^{j_1}, \dots, R_i^{j_{\beta(K_i)}}$ to obtain a more compact representation of the right filtration \mathcal{R} :

$$\begin{aligned} \mathcal{W}_i &= (W_i^1, \dots, W_i^{\beta(K_i)}) \\ &= (R_i^{j_1}, \dots, R_i^{j_{\beta(K_i)}}) \end{aligned} \tag{8}$$

where the $R_i^{j_k}$ are those with $\dim(R_i^{j_k}/R_i^{j_k-1}) = 1$, therefore the quotient spaces W_i^j/W_i^{j-1} are all one-dimensional. We note that a basis $\{[w_i^j]\}_{j=1}^{\beta(K_i)}$ for V_i is *compatible* with the right filtration \mathcal{W}_i if there is one basis element in each quotient space:

$$[w_i^j] \in W_i^j/W_i^{j-1}$$

for $j = 1, \dots, \beta(K_i)$. We return to this concept in Section 3.2.

Additionally, let

$$\mathbf{b}_i^W = (b_i^{j_1}, b_i^{j_2}, \dots, b_i^{j_{\beta(K_i)}})$$

contain the birth times of the non-zero quotient spaces, which is the birth vector for \mathcal{W} . So \mathbf{b}_i^W is a subset of the birth vector \mathbf{b}_i for the full right filtration \mathcal{R} .

The zigzag persistence algorithm is implemented by determining whether a birth or a death is occurring with each simplex addition or deletion. The right filtration and birth vector are then updated accordingly, and when a death occurs, the quotient space R_i^j/R_i^{j-1} corresponding to it is determined, and the associated birth time b_i^j used to output the interval $[b_i^j, i]$.

While the output of intervals $\{[b_j, d_j] \mid j = 1, \dots, m\}$ is unique, there may be more than one way to choose homology classes corresponding each interval. In Section 3 we will propose a method for choosing a homology class (by choosing a specific representative cycle for it) for each interval at each time point in a way that is geometrically motivated, but still compatible with the right filtration.

3 Tracking representative cycles

3.1 Motivation

Our interest in choosing and tracking representative cycles over a sequence of spaces stems from analysis of coverage holes in time-varying sensor networks. The idea of using homological methods to study coverage in sensor networks was proposed by de Silva and Ghrist ([7], [6]), and the use of zigzag persistent homology allows some of these ideas to be employed in the dynamic network setting. The set of intervals output from the zigzag persistence algorithm describes the birth and death times of homological features, and these features do not necessarily correspond to individual coverage holes [1]. Ideally, we are interested in tracking coverage holes over time, but this is not possible in general, given the constraints on the limited geometric information available with the sensor network model. Instead, we try and obtain a ‘good’ representative cycle

for a hole as it appears in the network, and then propagate this cycle over time as best as possible. Below we describe in more detail the model for the sensor network (3.1.1), the representative cycles we would ideally like to obtain (3.1.3), and those that we are able to compute (3.1.4).

3.1.1 Homology for sensor networks

A network consists of a set of sensors, each at the center of an isotropic coverage disk of radius r . The union of the disks creates the coverage region for the entire network, and we are interested in making statements about coverage properties of this network, as the sensors are allowed to move over time. A communication graph is created by connecting any two sensors by an edge when they are less than distance $2r$ from one another, and the homology of the Rips complex of this graph is used to approximate the homology of the coverage region of the network. Figure 1 shows the coverage region (left), communication graph (center) and associated Rips complex (right) for a given sensor network. Note that a Rips complex is the maximal simplicial complex that can be built from a given graph, but since we are only interested in computing the first homology we only need to consider the 2-skeleton of the Rips complex. The Rips complex includes a 2-simplex defined by three sensors whenever their coverage disks have nonempty pairwise intersections, so if the disks have no triplet-wise intersection then a small hole may be present in the coverage region which is not detected by the Rips complex. See Figure 2 (left) for an example of this. For our purposes we designate such holes as too small to be of importance, and work with the homology as it is defined by the Rips complex. See [7] for an alternative approach, which allows false alarms (holes in the complex which do not exist in the coverage region), but is able to give coverage guarantees.




Figure 1: (Left) The coverage region for a sensor network. (Center) The communication graph. (Right) The associated Rips complex.

A final key result that we mention is by Chambers et al [5], who show that the first homology of the Rips complex (a combinatorial object) is the same as the first homology of the projection of the Rips complex onto the plane (this projection is referred to as the Rips shadow). The Rips shadow corresponding to the sensor network from Figure 1 is shown in Figure 2 (right). For a Rips complex K , we denote its shadow as K^S .

3.1.2 Zigzag persistence for dynamic networks

In a time-varying network (represented by a sequence of simplicial complexes K_{t_1}, \dots, K_{t_T}), homology classes may be tracked over time using zigzag persistent homology by mapping through the union complexes. So the sequence of simplicial complexes




Figure 2: (Left) A situation where three coverage disks each intersect pairwise, but there is no point where all three intersect. These three sensors will form a ‘filled-in’ triangle in the Rips complex, but there is a hole in the coverage region. (Right) The Rips shadow corresponding to the Rips complex in Figure 1.

$$\begin{array}{ccccccc}
 K_{t_1} & & K_{t_2} & & \dots & & K_{t_T} \\
 \swarrow & & \searrow & & & & \swarrow \\
 & (K_{t_1} \cup K_{t_2}) & & (K_{t_2} \cup K_{t_3}) & & & (K_{t_{T-1}} \cup K_{t_T})
 \end{array} \tag{9}$$

gives rise to the associated zigzag persistence module:

$$\begin{array}{ccccccc}
 H_1(K_{t_1}) & & H_1(K_{t_2}) & & \dots & & H_1(K_{t_T}) \\
 \swarrow & & \searrow & & & & \swarrow \\
 & H_1(K_{t_1} \cup K_{t_2}) & & H_1(K_{t_2} \cup K_{t_3}) & & & H_1(K_{t_{T-1}} \cup K_{t_T})
 \end{array}$$

For implementational and theoretical purposes the sequence in Equation 9 is broken down, with each forward map re-written as a series of single simplex additions, and each backward map as a series of single simplex deletions. This refinement induces the analogous refinement on the zigzag module.

3.1.3 Canonical basis

Given a compact region in the plane such as the Rips shadow K^S , there exists a ‘canonical basis’ for its first homology space, where each basis homology class surrounds a single hole. Consider $\overline{K^S}$, the complement of K^S in \mathbb{R}^2 , then the number of separate components in $\overline{K^S}$ (ignoring the infinite component) is equal to the number of holes in K^S (i.e. the rank of $H_1(K^S)$). This result is a specific case of the more general principle of Alexander Duality (see, for example Ch. 5 of [10]), which for a certain class of spaces, relates the k -th reduced homology of a space to the $n - q - 1$ -th cohomology of the complement of the space (where n is the embedding dimension). We do not go into details here, but the salient point is that a canonical basis exists for the first homology of a space in the plane, with one homology class surrounding each hole.

Since the Rips complex K and its Rips shadow K^S have the same homology, a desirable goal would be to have a homology basis for the Rips complex, where projection of this basis onto the Rips shadow gives the canonical basis. In particular, we would like a representative cycle for each homology class in the basis, where the projection of the representative cycle onto the Rips shadow is homologous to the boundary of one of the holes. *In general, this desirable goal is not possible.* The Rips complex itself is not embeddable in two dimensions, so Alexander Duality cannot be applied to obtain a canonical basis for its first homology. Moreover, although K has the same homology as

K^S , it is impossible to know whether a given homology basis for K corresponds to the canonical basis or not (without knowing coordinates for the vertices, or the projection map from K onto K^S).

We will see in Section 3.1.4 that taking the dynamic nature of the network into account, there are some cases where it is possible to make a canonical choice for a homology class (with corresponding representative cycle) at its birth or death time. In Section 3.2 we present a method for obtaining these cycles and updating them as the network changes over time, along with an explicit algorithm for doing so.

3.1.4 Partial canonical information

As described in Table 1, a homology class can be born by either the addition or removal of a simplex, and similarly a death is caused by either the addition or removal of a simplex, resulting in four distinct cases for how the homology can change. In this section we illustrate the two cases corresponding to ‘births’, and how one of them allows a canonical choice of homology class. Our discussions here are with respect to the first homology, but the same principles hold for d -dimensional homology.

In the sequence of simplicial complexes, the birth of a homology class occurs at time i when either the forward map $V_i \xrightarrow{f_i} V_{i+1}$ has nonzero cokernel, or the backward map $V_i \xleftarrow{g_i} V_{i+1}$ has nonzero kernel. Of these two cases, $\ker(g_i) \neq \mathbf{0}$ is the only one which indicates the specific homology class that is being born.

Consider the case where the birth is in first homology. If a hole is formed by the removal of a 2-simplex, then there is a unique homology class (the one surrounding the hole) which is born. This homology class also the unique homology class in $\ker(g_i)$ (i.e.: the only homology class that is nontrivial in K_{i+1} but trivial in K_i). On the other hand, if a hole is formed by the addition of an edge, there are many choices for which homology class is being born, with no choice being canonical. For example, if a hole is split into two by the addition of an edge, then which of them is the ‘new’ hole? See the first two rows of Figure 3 for an illustration of these cases.

Our approach then, is to maintain a basis for the homology at each time point, making the canonical choice of homology class whenever available, and tracking that choice through the sequence of complexes as best as possible. Our method for implementing this, along with the specific basis we maintain and its relation to the zigzag algorithm, is detailed in Section 3.2.

3.2 Algorithm

As mentioned in Section 4, a zigzag module \mathbb{V} (Equation 1) has unique interval decomposition $\text{Pers}(\mathbb{V}) = \{[b_j, d_j] \mid j = 1, \dots, m\}$, which describes the birth and death times of homological features in the sequence. This decomposition is determined through the maintenance of a right filtration \mathcal{W}_i (Equation 8) on the space V_i , and a birth vector \mathbf{b}_i for $i = 1, \dots, n$. The zigzag persistence algorithm performs this task by determining whether a birth or a death is occurring for each simplex σ being added or removed, and updating the right filtration and birth vector accordingly (and outputting the appropriate birth-death interval whenever a death occurs).

At each stage in our algorithm, we maintain a basis for the homology that attempts to approximate the canonical basis as best as possible. Further, the basis homology classes are compatible with the right filtration \mathcal{W}_i , in the sense that the j^{th} basis homology class is an element of the j^{th} quotient space W_i^j / W_i^{j-1} . This means that the span of the first j homology classes in the basis is equal to the j^{th} subspace W_i^j in the right filtration \mathcal{W}_i . This property is necessary if we wish to interpret our basis homology

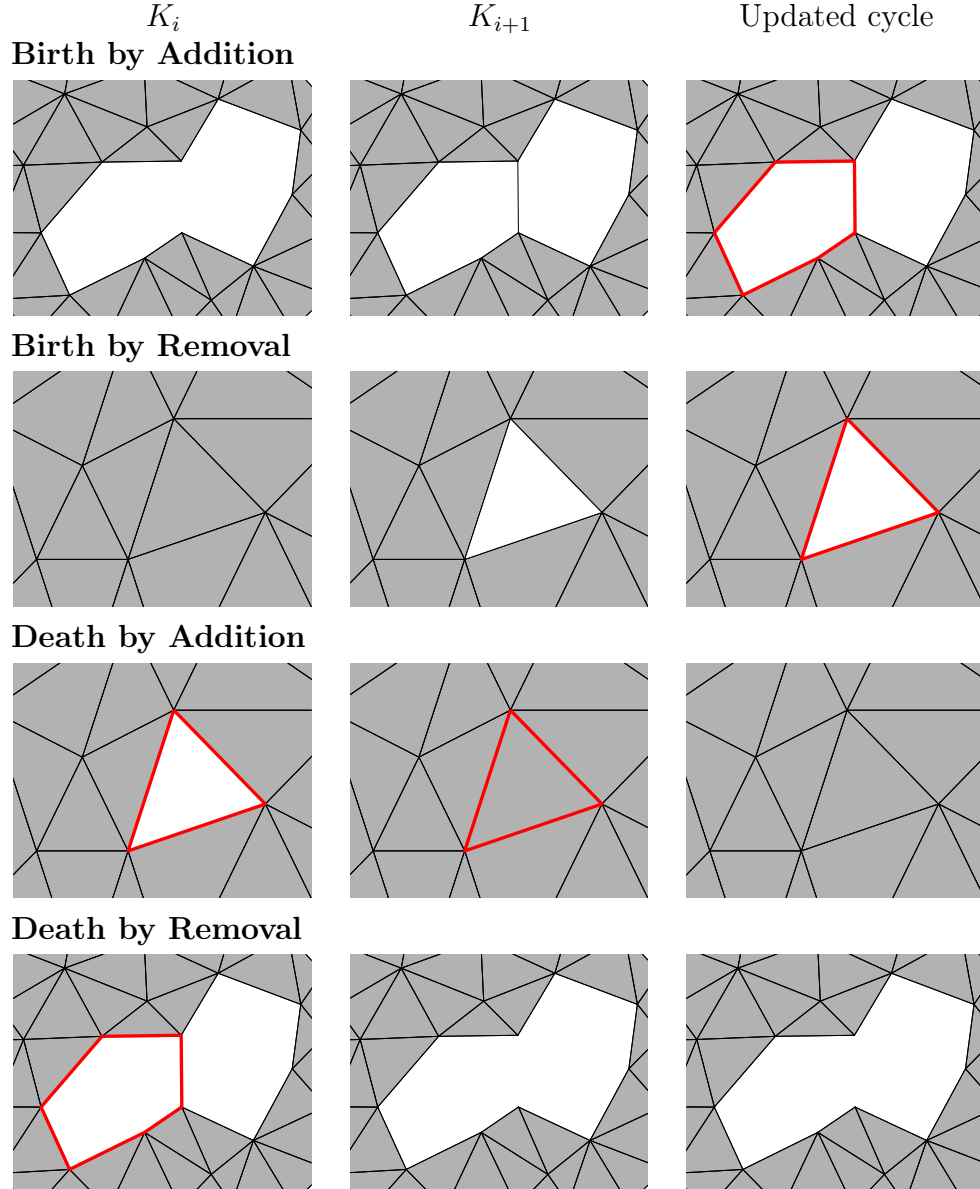


Figure 3: Simplicial complexes which illustrate the four cases where the first homology changes, and the corresponding changes to the representative cycles.

classes as corresponding to particular intervals in the birth-death decomposition. The intervals are really describing specific quotient spaces in the right filtration that have persisted over the sequence, so our homology classes need be assigned one-to-one to the quotient spaces. The proof that this property is maintained during the algorithm is presented in Section 3.2.

The method we present here is computed using the regular zigzag persistent homology algorithm, but keeps an explicit record of the homology basis chosen for the right filtration at each time point. The homology basis is stored by choosing a specific representative cycle for each basis homology class. The choice of basis homology classes is not unique, so it is made in a geometrically meaningful way, attempting to approximate the canonical basis. The zigzag persistence algorithm supplies information about whether the addition or removal of simplex σ is causing a birth or a death. If it is a birth, σ is called a positive simplex, denoted σ^+ , and if it is a death, σ is called a negative simplex, denoted σ^- .

When a birth occurs, we must add a new representative cycle to our list. As mentioned in Section 3.1.4, when a birth occurs due to the removal of a simplex σ , there is a canonical choice available for the new homology class. We choose the boundary $\partial\sigma$ of the removed simplex as the representative cycle for this homology class, since it is the shortest cycle surrounding the new hole. When the birth occurs due to the addition of a simplex σ , there is no canonical choice for which is the ‘new’ homology class, but any cycle containing σ will have its homology class in $\text{coker}(f_i)$. As a heuristic, we choose the shortest cycle containing σ as the new representative cycle.

When a death occurs, we must remove a representative cycle from our list. Analogous to the two ways in which a birth can occur, a death occurs when either the forward map $V_i \xrightarrow{f_i} V_{i+1}$ has nonzero kernel, or the backward map $V_i \xleftarrow{g_i} V_{i+1}$ has nonzero cokernel. Of these two cases, $\ker(f_i) \neq \mathbf{0}$ is the only one which indicates the specific homology class $[c] = \ker(f_i)$ that is being killed (becoming trivial). In this case, the death occurs due to the addition of a simplex, and we reduce the matrix storing the representative cycles with respect to the boundary matrix ∂ , and remove the cycle which becomes trivial. If the death occurs due to the removal of a simplex σ , then the first representative cycle containing σ is removed, and a change of basis is performed to remove σ from any remaining representative cycles. This is done in the same way as the change of basis operation in the regular zigzag persistence algorithm.

The four cases are illustrated in Figure 3. It is worth noting the distinction between the two types of deaths. In a death by addition, the representative cycle is still present in the complex, but becomes homologous to zero. In a death by removal, the representative cycle whose homology class is getting killed is no longer present in the complex, and it is possible that other representative cycles are also no longer present in the complex, so another representative cycle which is from the same quotient space must be chosen.

We store the representative cycles for time i in the matrix W_i , which is retained for all time points. The algorithm is summarized below.

```

%Notation
 $w_i^l$  = column  $l$  of  $W_i$ 
 $w_i^j[\sigma]$  = coefficient of  $\sigma$  in  $w_i^j$ 
 $\partial_d$  = boundary matrix

%Initialize
 $W_0 = n \times 0$  matrix
 $b_0$  = empty vector

%Perform updates
for  $i = 1$  to  $n$ 
    if  $K_i = K_{i-1} - \{\sigma\}$                                 %simplex removal
        if  $\sigma^+$                                           %birth
             $W_i = [\partial\sigma_i \ W_{i-1}]$                 %prepend  $\partial\sigma$ 
             $b_i = [b_{i-1} \ i]$ 
        else if  $\sigma^-$                                       %death
             $l =$  index of first nonzero entry in  $r_\sigma$ 
             $c_l = r_\sigma(l)$ , the coefficient of  $\sigma$  in  $w_l$ 
            for  $j = 1$  to (# columns of  $W_{i-1}$ )                %change of basis
                 $c_j = r_\sigma(j)$ , the coefficient of  $\sigma$  in  $w_j$ 
                 $w_j = w_j - \frac{c_j}{c_l} w_l$ 
            end
             $W_i = W_{i-1}$  with column  $l$  and row  $r_\sigma$  removed
             $b_i = b_{i-1}$  with entry  $l$  removed
        end
    end
    if  $K_{i+1} = K_i \cup \sigma$                                 %simplex addition
        if  $\sigma^+$                                           %birth
             $W_i = [W_{i-1} \ (Cu - \sigma)]$                 %append  $Cu - \sigma$ 
             $b_i = [b_{i-1} \ i]$ 
        else if  $\sigma^-$                                       %death
             $l =$  index for col of  $W_{i-1}$  trivial when  $[\partial_d \ W_{i-1}]$  reduced
             $W_i = W_{i+1}$  with column  $w_l$  removed
        end
    end
end
end

```

4 Correctness

The adaptive representative cycles that we track using the algorithm described in Section 3.2 are stored as column vectors w_i^k in a matrix \mathbf{W}_i

$$\mathbf{W}_i = [w_i^1 \ w_i^2 \ \dots \ w_i^{\beta(K_i)}]$$

In this section we show that the homology classes represented by the cycles w_i^k form a basis for V_i , and further that their order in \mathbf{W}_i corresponds to the order of the right filtration \mathcal{W}_i (Equation 8). In other words, the span of the homology classes of the first j representative cycles is equal to the j^{th} space W_i^k in the filtration \mathcal{W}_i of V_i .

i.e.:

$$\text{span}\{[w_i^k]\}_{k=1}^j = W_i^j \tag{10}$$

for $j = 1, \dots, \beta(K_i)$.

We prove this by induction on i , beginning with the base case of a single vector space $\mathbb{V} = V_1$, which results from a simplicial complex of one vertex $K_1 = \sigma$. This yields

$$\begin{aligned}\mathcal{W}_1 &= (V_1) \\ \mathbf{W}_1 &= [w_1^1]\end{aligned}$$

where $w_1^1 = [1]$ is the column vector of length 1 representing the cycle consisting of the vertex σ . The homology class $[w_1^1]$ spans the one-dimensional homology space $W_1^1 = V_1$.

In the inductive step, we assume that for

$$\mathcal{W}_i = (W_i^1, \dots, W_i^{\beta(K_i)}) \quad (11)$$

$$\mathbf{W}_i = [w_i^1 \ w_i^2 \ \dots \ w_i^{\beta(K_i)}] \quad (12)$$

we have (for $j = 1, \dots, \beta(K_i)$)

$$\text{span}\{[w_i^k]\}_{k=1}^j = W_i^j$$

We will show then that (for $j = 1, \dots, \beta(K_{i+1})$)

$$\text{span}\{[w_{i+1}^k]\}_{k=1}^j = W_{i+1}^j \quad (13)$$

for all four of the cases described in the algorithm (Section 3.2). In all cases we use σ to denote the d -simplex being added or removed, and the updates are performed on the representative cycles and right filtration of the appropriate dimension (see Table 1).

1. Birth by addition. The map $V_i \xrightarrow{f_i} V_{i+1}$ has $\text{coker}(f_i) \neq \mathbf{0}$, and the new right filtration is

$$\mathcal{W}_{i+1} = (f_i(W_i^0), f_i(W_i^1), \dots, f_i(W_i^{\beta(K_i)}), V_{i+1})$$

where $V_{i+1}/f_i(W_i^{\beta(K_i)}) = \text{coker}(f_i)$.

The new list of representative cycles is

$$\mathbf{W}_{i+1} = [\mathbf{W}_i^{\sigma^+} \ w_{new}]$$

where $\mathbf{W}_i^{\sigma^+}$ is the matrix \mathbf{W}_i with an additional row of zeros added, corresponding to simplex σ (so the cycles are now written in terms of simplices of K_{i+1} instead of simplices of K_i), and w_{new} is a cycle in K_{i+1} containing σ . There is no canonical choice for which cycle containing σ should be chosen, and our proof holds regardless of the choice. As mentioned in Section 3.2, we make this choice based on shortest hop length.

Since $w_{i+1}^k = w_i^k$ as chains (with the appropriate row for σ added containing a 0 coefficient), we get $[w_{i+1}^k] = f_i([w_i^k])$, for $k = 1, \dots, \beta(K_i)$ because f_i is the map induced by inclusion. Therefore

$$\begin{aligned}W_{i+1}^j &= f_i(W_i^j) \\ &= f_i\left(\text{span}\{[w_i^k]\}_{k=1}^j\right) \\ &= \text{span}\{f_i([w_i^k])\}_{k=1}^j \\ &= \text{span}\{[w_{i+1}^k]\}_{k=1}^j\end{aligned}$$

for $j = 1, \dots, \beta(K_i)$.

Finally, we must show that $[w_{new}]$ is nontrivial and in $\text{coker}(f_i)$, and therefore linearly independent from $\{[w_i^j]\}_{j=1}^{\beta(K_i)}$, so they together span the $\beta(K_1) + 1 = \beta(K_{i+1})$ -dimensional vector space $V_{i+1} = W_{i+1}^{\beta(K_{i+1})}$. First, note that having a nonzero coefficient for σ in w_{new} : that $[w_{new}] \neq 0$; and that any cycle c in the same homology class $[w_{new}]$ will also have a nonzero coefficient for σ . These are due to the fact that σ is not contained in the boundary of any other simplex, and the difference between homologous cycles must be written as a linear combination of boundaries (therefore the coefficient for σ is zero in the difference $c - w_{new}$, but is nonzero in w_{new} , so must also be nonzero in c). Moreover, note that $[w_{new}] \not\subseteq \text{im}(f_i)$, since any homology class in $\text{im}(f_i)$ must have a representative cycle in the image under inclusion $i(K_i) \subset K_{i+1}$, and all cycles in $[w_{new}]$ contain $\sigma \notin i(K_i)$. Therefore, we have

$$\begin{aligned} W_{i+1}^{\beta(K_{i+1})} &= V_i \\ &= \text{im}(f_i) \oplus \text{coker}(f_i) \\ &= \left(\text{span}\{[w_{i+1}^k]\}_{k=1}^{\beta(K_i)} \right) \oplus [w_{new}] \\ &= \text{span}\{[w_{i+1}^k]\}_{k=1}^{\beta(K_{i+1})}. \end{aligned}$$

as desired.

2. Birth by removal. The map $V_i \xleftarrow{g_i} V_{i+1}$ has $\ker(g_i) \neq \mathbf{0}$, and the new right filtration is

$$\mathcal{W}_{i+1} = (\ker(g_i), g_i^{-1}(W_i^1), g_i^{-1}(W_i^2), \dots, g_i^{-1}(W_i^{\beta(K_i)})).$$

This is because in the full right filtration

$$\mathcal{R}_{i+1} = (\mathbf{0}, g_i^{-1}(R_i^0), g_i^{-1}(R_i^1), \dots, g_i^{-1}(R_i^i))$$

if R_i^j / R_i^{j-1} was nontrivial in \mathcal{R}_i , then $g_i^{-1}(R_i^j) / g_i^{-1}(R_i^{j-1})$ will be nontrivial in \mathcal{R}_{i+1} for $j = 1, \dots, i$. This means that if W_i^j is a subspace in \mathcal{W}_i then $g_i^{-1}(W_i^j)$ is a subspace in \mathcal{W}_{i+1} . Also, the new term

$$g_i^{-1}(R_i^0) / \mathbf{0} = g_i^{-1}(\mathbf{0}) / \mathbf{0} = \ker(g_i) / \mathbf{0} = \ker(g_i)$$

is nontrivial, and yields the first term $\ker(g_i)$ in \mathcal{W}_{i+1} .

The new list of representative cycles is

$$\mathbf{W}_{i+1} = [\partial\sigma \ \mathbf{W}_i]$$

where $\partial\sigma$ are the simplices that make up the boundary of σ , but considered in K_{i+1} , instead of K_i .

First we note that $\text{span}\{[\partial\sigma]\} = \ker(g_i)$. This is because $i(\partial\sigma)$ is the boundary of simplex σ in K_i and hence homologous to zero, thus

$$g_i([\partial\sigma]_{i+1}) = [\partial\sigma]_i = \mathbf{0}$$

which means $[\partial\sigma] \subseteq \ker(g_i)$. The cycle $\partial\sigma$ is also nontrivial in K_{i+1} , because if there exists a d -chain c in K_{i+1} that had $\partial\sigma$ as its boundary, then in K_i the union of σ with $i(c)$ in K_i would form a d -cycle, and the removal of σ would result in the death of that d -cycle, instead of the birth of a $(d-1)$ -cycle, which is a contradiction. Therefore, $[\partial\sigma]$ spans a one-dimensional subspace of the one-dimensional space $\ker(g_i)$, so $\text{span}\{[\partial\sigma]\} = \ker(g_i)$.

Now we show that $W_{i+1}^j = \text{span}\{[w_{i+1}^k]\}_{k=1}^j$ for $j = 1, \dots, \beta(K_{i+1})$. First note the index change, so

$$w_{i+1}^{k+1} = w_i^k$$

for $k = 1, \dots, \beta(K_i)$. Consider the representative cycle w_i^k , and another cycle c which is homologous to w_i^k in K_i . Since c and w_i^k are both $(d-1)$ -cycles, they are also present in K_{i+1} . Then $[c]_i = [w_i^k]_i$ implies $[c]_{i+1} = [w_i^k]_{i+1} + a[\partial\sigma]_{i+1}$, where $a = 0$ or 1 . Therefore

$$g_i^{-1}([w_i^k]) = [w_{i+1}^{k+1}] \oplus [\partial\sigma].$$

So

$$\begin{aligned} W_{i+1}^j &= g_i^{-1}(W_i^{j-1}) \\ &= \text{span}\{g_i^{-1}[w_i^k]\}_{k=1}^{j-1} \\ &= \text{span}\left\{[w_{i+1}^{k+1}] \oplus [\partial\sigma]\right\}_{k=1}^{j-1} \end{aligned}$$

for $j = 2, \dots, \beta(K_{i+1})$. Combining this with

$$W_{i+1}^1 = \ker(g_i) = [\partial\sigma] = [w_{i+1}^1]$$

we obtain

$$W_{i+1}^j = \text{span}\{[w_{i+1}^k]\}_{k=1}^j$$

for $j = 1 \dots, \beta(K_{i+1})$, as desired.

3. Death by addition. For the map $f_i : V_i \rightarrow V_{i+1}$ we get $\ker(f_i) = [\partial\sigma]$ with a similar proof to that of case **2** above.

Since $\ker(f_i) \neq \mathbf{0}$, we have $\text{coker}(f_i) = \mathbf{0}$, so $V_{i+1}/f_i(V_i) = \mathbf{0}$. Also, there exists an index $l \in \{1, \dots, \beta(K_i)\}$ such that $[\partial\sigma] \in W_i^l$, but $[\partial\sigma] \notin W_i^{l-1}$ (using the convention $W_i^0 = \mathbf{0}$), so

$$f_i(W_i^l/W_i^{l-1}) = \mathbf{0}$$

This gives

$$W_{i+1} = (f_i(W_i^1), \dots, f_i(W_i^{l-1}), f_i(W_i^{l+1}), \dots, f_i(W_i^{\beta(K_i)}))$$

so we have

$$W_{i+1}^j = \begin{cases} f_i(W_i^j) & \text{if } j < l; \\ f_i(W_i^{j+1}) & \text{if } j \geq l. \end{cases} \quad (14)$$

Considering now the representative cycles, we need to determine the index l . Since the elements $\{[w_i^k]\}_{k=1}^{\beta(K_i)}$ form a basis for V_i , we can write uniquely

$$[\partial\sigma] = \sum_{k=1}^{\beta(K_i)} \alpha_k [w_i^k]. \quad (15)$$

Then $[\partial\sigma] \in \text{span}\{[w_i^k]\}_{k=1}^l = W_i^l$, but $[\partial\sigma] \notin \text{span}\{[w_i^k]\}_{k=1}^{l-1} = W_i^{l-1}$ implies that α_l is the last nonzero coefficient in this sum. We now define

$$w_{\text{remove}} = w_i^l$$

and obtain

$$\mathbf{W}_{i+1} = [w_i^1 \ \dots \ w_i^{l-1} \ w_i^{l+1} \ \dots \ w_i^{\beta(K_i)}]$$

noting that all of the simplices in the $(d-1)$ -cycles w_i^k are present in K_{i+1} . Therefore the corresponding homology classes are related by

$$[w_{i+1}^j] = \begin{cases} f_i([w_i^j]) & \text{if } j < l; \\ f_i([w_i^{j+1}]) & \text{if } j \geq l. \end{cases}$$

for $j = 1, \dots, \beta(K_{i+1})$, since f_i is the map induced by inclusion. This, together with Equation 14 yields

$$W_{i+1}^j = \text{span}\{[w_{i+1}^k]\}_{k=1}^j$$

for $j = 1, \dots, \beta(K_{i+1})$, as desired.

Note that the index l indicating the last nonzero coefficient in Equation 15 also determines the birth-death interval: $[\mathbf{b}_i^W[l], i]$.

4. Death by removal. The map $V_i \xleftarrow{g_i} V_{i+1}$ has $\text{coker}(g_i) \neq \mathbf{0}$. There exists an index l such that $W_i^j \subseteq \text{im}(g_i)$, for all $j < l$, but $W_i^l \not\subseteq \text{im}(g_i)$. Then

$$g_i^{-1}(W_i^l / W_i^{l-1}) = \mathbf{0}$$

so

$$\mathcal{W}_{i+1} = (g_i^{-1}(W_i^1), \dots, g_i^{-1}(W_i^{l-1}), g_i^{-1}(W_i^{l+1}), \dots, g_i^{-1}(W_i^{\beta(K_i)})).$$

We note that the image of this in V_i is

$$\begin{aligned} g_i(\mathcal{W}_{i+1}) &= \mathcal{W}_i / \text{coker}(g_i) \\ &= (W_i^1, \dots, W_i^{l-1}, W_i^{l+1} / \text{coker}(g_i), \dots, W_i^{\beta(K_i)} / \text{coker}(g_i)) \end{aligned} \quad (16)$$

Considering now the representative cycles, l is the index of the first representative cycle w_i^l which contains σ . To see that this is the same index l as described above, note that since w_i^k doesn't contain σ for $k < l$, we have $[w_i^k] \in \text{im}(g_i)$, and $\text{span}\{[w_i^k]\}_{k=1}^j = W_i^j \subseteq \text{im}(g_i)$, for all $j < l$, but $\text{span}\{[w_i^k]\}_{k=1}^l = W_i^l \not\subseteq \text{im}(g_i)$.

Denoting the coefficient for σ in representative cycle w_i^k by $w_i^k[\sigma]$, we consider another set of representative cycles in K_i

$$\hat{w}_i^k = w_i^k - \frac{w_i^k[\sigma]}{w_i^l[\sigma]} w_i^l.$$

By definition, σ is not present in any \hat{w}_i^k , so we are able to define

$$w_{i+1}^k = \begin{cases} \hat{w}_i^k & \text{if } k < l; \\ \hat{w}_i^{k+1} & \text{if } k \geq l. \end{cases}$$

to be our representative cycles in K_{i+1} , with the row corresponding to σ removed. Then

$$\begin{aligned} \mathbf{W}_{i+1} &= [\hat{w}_i^1 \ \dots \ \hat{w}_i^{l-1} \ \hat{w}_i^{l+1} \ \dots \ \hat{w}_i^{\beta(K_i)}] \\ &= [w_{i+1}^1 \ \dots \ w_{i+1}^{l-1} \ w_{i+1}^l \ \dots \ w_{i+1}^{\beta(K_{i+1})}] \end{aligned}$$

We proceed by showing that the \hat{w}_i^k completely determine the quotiented filtration $\mathcal{W}_i / \text{coker}(g_i)$ in Equation 16, in the sense that

$$W_i^j / \text{coker}(g_i) = \text{span}\{[\hat{w}_i^k]\}_{k=1}^j \quad (17)$$

for $j = 1, \dots, \beta(K_{i+1})$.

To show that Equation 17 holds, we show it separately for $j < l$, $j = l$, and $j > l$. For the first case, note that when w_i^k does not contain σ , we have $\hat{w}_i^k = w_i^k$. In particular, for $k < l$ we have $\hat{w}_i^k = w_i^k$, therefore

$$W_i^j / \text{coker}(g_i) = W_i^j = \text{span}\{[w_i^k]\}_{k=1}^j = \text{span}\{[\hat{w}_i^k]\}_{k=1}^j$$

for $j = 1, \dots, l-1$.

By assumption $W_i^l \not\subseteq \text{im}(g_i)$, but $W_i^{l-1} \subseteq \text{im}(g_i)$, so

$$W_i^l / \text{coker}(g_i) = W_i^{l-1} = \text{span}\{[\hat{w}_i^k]\}_{k=1}^{l-1} = \text{span}\{[\hat{w}_i^k]\}_{k=1}^l$$

since $\hat{w}_i^l = \vec{0}$.

For $j > l$, we first note that the homology elements $\{[\hat{w}_i^k]\}$ are linearly independent for $k \in \{1, \dots, l-1, l+1, \dots, \beta(K_i)\}$. This is because each $[\hat{w}_i^k]$ is a subset of $[w_i^k] \oplus [w_i^l]$ (but not equal to $[w_i^l]$), and the $\{[w_i^k]\}$ are linearly independent. So $\{[\hat{w}_i^k]\}_{k=1}^j$ span a $(j-1)$ -dimensional space when $j > l$ (since $[w_i^l]$ is trivial). Also, because all the \hat{w}_i^k have a zero coefficient for σ , they are not in the $\text{coker}(g_i)$. So

$$\text{span}\{[\hat{w}_i^k]\}_{k=1}^j \subseteq W_i^j / \text{coker}(g_i)$$

Moreover, we note that $W_i^j / \text{coker}(g_i)$ is also a $(j-1)$ -dimensional space for $j > l$. So $\text{span}\{[\hat{w}_i^k]\}_{k=1}^j = W_i^j / \text{coker}(g_i)$.

It now follows that since

$$g_i([w_{i+1}^k]) = \begin{cases} [\hat{w}_i^k] & \text{if } k < l; \\ [\hat{w}_i^{k+1}] & \text{if } k \geq l \end{cases}$$

then

$$\begin{aligned} W_{i+1}^j &= \begin{cases} g_i^{-1}(W_i^j) = g_i^{-1}(W_i^j / \text{coker}(g_i)) & \text{if } j < l; \\ g_i^{-1}(W_i^{j+1}) = g_i^{-1}(W_i^{j+1} / \text{coker}(g_i)) & \text{if } j \geq l \end{cases} \\ &= \text{span}\{[w_{i+1}^j]\}_{i+1}^j \end{aligned}$$

completing the induction.

5 Applications to dynamic sensor networks

In this Section, we illustrate some of the types of results possible when using the representative cycles to ‘track’ coverage holes in dynamic sensor networks. A more detailed description of these applications is available in [9].

5.1 Tracking coverage holes

A typical model for a dynamic sensor network has n nodes distributed randomly uniformly over a region of interest (say the unit square, for simulation purposes). The nodes are then allowed to move independently and randomly according to the same stochastic process, such as a discrete-time Brownian motion. The coverage region and Rips complex for this network are described in Section 3.1.1, and we note that the only information necessary to build the Rips complex is the binary adjacency matrix indicating which nodes are within distance $2r$ of each other. The input to the algorithm then, is simply a sequence of adjacency matrices, representing the communication graph of the network at each time point. No information about the locations of the nodes, or distances between them is used. The output is a barcode, along with a representative cycle for each bar at each time point. These representative cycles can be used to track holes in many circumstances.

Since canonical information is only available when a hole is formed due to the removal of a triangle, this method performs best when the network is relatively dense, and the holes are appearing and disappearing in relative isolation (without extensive merging and splitting of holes during their lifetimes). In such a setting the representative cycles typically surround each hole uniquely, and even encircle the holes relatively tightly. Figure 4 illustrates this, with the color of each cycle indicating the bar that it corresponds to in the barcode.




Figure 4: Representative cycles, color-coded with their corresponding intervals (shown in the barcode - top left).

Hop-length of shortest cycle surrounding hole	Persistence of hole in hop distance filtration
4, 5, 6	1
7, 8, 9	2
10, 11, 12	3
\vdots	\vdots
$3k + 1, 3k + 2, 3k + 3$	k

Table 2: The relationship between persistence of a cycle in the hop-distance filtration, and size (in terms of hop-length) of the largest hole in surrounds.

5.2 Enhancing barcode with estimated size information

In addition to tracking holes, choosing a specific representative cycle for each bar allows us to attach additional information onto the barcode, such as size estimates at each point. Given our coordinate-free setting, the only size estimate available for a hole is the length of the shortest cycle surrounding it. For a specific simplicial complex K , this may be obtained by performing a hop-distance based filtration. Such a filtration starts with the original simplicial complex $K^1 = K$, and then forms a nested sequence of complexes K^2, K^3, \dots , where K^h is formed by taking K^{h-1} and adding an edge between nodes that are h hops apart in K , and then filling in all higher-dimensional simplices. A hop-distance filtration on a simple simplicial complex consisting of a single loop is shown in Figure 5. Given a nontrivial cycle in K , the depth at which it becomes trivial in the hop-distance filtration is determined by the size (in shortest hop-length) of the largest hole that the cycle surrounds. Table 2 gives this relationship.

Now, a hop-distance filtration is performed on the simplicial complex at each time point, and for each bar in the zigzag barcode, the persistence of its representative cycle in the hop distance filtration is computed. This additional size information may be added to the barcode visually by thickening the bar proportionally to the size value.

As an example, Figure 6 shows a network with an expanding failure region. The barcode thickened by the estimated hole size is shown in Figure 7, and it is clear that there is one hole that is becoming increasingly problematic.




Figure 5: Hop distance filtration at depth $k = 1$ (original complex), left, $k = 2$ (hole persists), middle, and $k = 3$ (hole is filled in), right.




Figure 6: A network with an expanding failure region, where nodes within the failure region can no longer sense or communicate.




Figure 7: Barcode from zigzag persistence on the network with the expanding failure region in Figure 6. The thickness of each bar is proportional to the depth its adaptive representative cycle persists in the hop distance filtration at each time point.

6 Conclusion

Persistent homology and zigzag persistent homology represent the dynamic homology of a sequence of spaces by computing a set of intervals, describing the birth and death times of homological features in the sequence. We present here a method for assigning a representative cycle at each time point to each interval in this decomposition. The original choice and method for updating these representative cycles are geometrically motivated, so they are interpreted as ‘tracking’ homological features. We describe the properties that the set of representative cycles must have to be compatible with the birth-death decomposition, and prove that the method presented maintains such properties.

Some applications of the method to track coverage holes in time-varying sensor networks are presented. For spaces in the plane, this method of tracking attempts to approximate the canonical basis for the first homology (where one homology class surrounds each hole), as best as possible, while still being compatible with the birth-death decomposition. Having chosen a specific representative cycle for each interval at each time point, additional features (such as estimates of hole size) can be attached onto the barcode, for a more comprehensive description of the dynamic coverage of the network.

References

- [1] Henry Adams and Gunnar Carlsson. Evasion paths in mobile sensor networks. *arXiv preprint arXiv:1308.3536*, 2013.
- [2] G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46:255–308, 2009.
- [3] G. Carlsson, V. de Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. *Proc. 25th Annual Symposium on Computational Geometry (SoCG)*, pages 247–256, June 2009.
- [4] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.

- [5] Erin W Chambers, Vin De Silva, Jeff Erickson, and Robert Ghrist. Vietoris–rips complexes of planar point sets. *Discrete & Computational Geometry*, 44(1):75–90, 2010.
- [6] V. de Silva and R. Ghrist. Homological sensor networks. *Notices of the American Mathematical Society*, 54(1):10–17, Jan 2007.
- [7] Vin de Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *The International Journal of Robotics Research*, 25(12):1205–1222, 2006.
- [8] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 1st edition, 2001.
- [9] H. Chintakunta J. Gamble and H. Krim. Coordinate-free quantification of coverage in dynamic sensor networks. Preprint, 2014.
- [10] Ezra Miller and Bernd Sturmfels. *Combinatorial commutative algebra*, volume 227. Springer, 2005.